Pre-viz Technologies for Standard Production Workflows

Technicolor Research and Innovation

1 Pre-viz challenges in CGI productions

CGI animated films and VFX relie on massively detailed 3D scenes featuring *hundreds of lights*, high resolution textures and many complex shaders. Pre-visualization of such assets involve complex computations, usually requiring the use of production renderers such as MentalRay[1] or RenderMan[2]. However, the non-interactive render times (up to several hours per frame) preclude any interactive setup of scene contents and lighting. Graphics hardware (GPU) accelerated solutions exists, they are generally based on progressive ray-tracing methods and thus very accurate to render shadows and reflections. However, they generally require the author to model the scenes in a specific way, for instance using proprietary shaders or lights, in order to get proper pre-viz. In this way they are not compatible with a standard production workflow based on Maya[3]/MentalRay. Other proprietary solutions exists but they are generally tightly bound to large CGI companies workflows and hardly applicable to other pipelines.

2 Technicolor's pre-viz technologies

In this talk, we present Technicolor's GPU pre-viz system compatible with standard Maya/MentalRay production workflows. This GPU-based pre-viz system has been developed upon Technicolor's research work in the field of real time rendering and pre-viz technology. Our solution addresses some of the common issues and bottlenecks found in Animation and VFX production workflows. It has been used in several productions such as *Prometheus* feature film or *Barbie with pink shoes* DVD. The following sections highlight the key features of our system and related research works that will be addressed during this talk.

Many-core animation system Our solution integrates a Many-Core Event Evaluation framework [Marvie et al. 2013] for real-time execution of many complex animation schemes. This technique takes advantages of task parallelism on many-core CPU architecture using a two-level scheduling approach. Our generic approach can deal with tens of thousands event-processing nodes, event loops, non-deterministic and interaction-driven animation modification at runtime.

High resolution texture management The use of high definition images (4K to 8K floating point encoded) in CGI production leads to very high memory consumptions hardly handled by standard GPU renderers. Building upon the ideas of low-redundancy and chunk reads, we developed an approach based on a tiled pyramidal format described in [Marvie and Bouatouch 2003]. This format can be seen as a progressive PNG format dedicated to texture maps. It provides an efficient texture representation allowing for dynamic loads and unloads of mipmap levels. In order to minimize the network fetches and transfers we developed an automatic differential caching system, accounting for remote texture file modifications. Such a mechanism minimizes the time to wait for HD texture maps to be available in the viewport pre-viz, thus providing a rich collaborative authoring experience.

Many lights, many shadows On the lighting side, most productions do not rely only on a small set of light emitters combined with a complex global illumination solution. Most of lighting is tweaked with fill- and rim- lights for artistic and performance reasons, generally leading to an important number of light sources in scenes. Contrary to existing OpenGL / DirectX-based viewports, our system supports interactive to real-time rendering with hundreds of dynamic lights at once. Through the design of a batching system that minimizes draw calls, dozens of shadow maps can also be rendered at interactive frame rates. This allows artists to setup their whole lighting set interactively and at once, without intermediate time-consuming renders.

Advanced lighting The use of reflectors at lighting stage is a common practice to highlight characters, reduce contrast or perform light balancing to match some artistic criteria like romantic atmosphere. Area lights fulfill these requirements and are widely used by lighting artist to perform such adjustment. Area lighting usually requires discrete evaluation of a surface integral generally achieved by a costly point light source sampling. We implemented an analytic solution that support both diffuse and specular area lighting, outperforming sampling based approaches.

Massive volumetric effects Scenes featuring many complex participating media are common in the world of production rendering. However, the complex inner structure of participating media requires vast amounts of memory for storage and costly computations. Such media are usually generated independently and arranged using crude representations. Technicolor's GPU pre-viz system implements a simple yet efficient out-of-core method for rendering massive volumetric scenes using Boundary-Aware Extinction Maps [Gautron et al. 2013]. Our solution allows interactive editing and rendering of massive heterogeneous volumetric datasets, providing to VFX artists a high valuable tool for the previsualization and tuning of complex volumetric environments for production rendering [Esteve et al. 2012].

Post-processing In addition to standard rendering, our solution integrates some advanced real-time post-processing techniques. Traditionally, during modeling and animation setup, artists do not have lighting information yet. As such, without proper shadowing, aligning correctly 3D objects while avoiding undesired penetrations or lack of contacts remains a tedious task. We developed a post-processing tool used to assess objects contacts validity, based on a bilateral filter that accounts for objects identifiers instead of the traditional luminance range. This is a time-saving feature that avoids most back-and-forth corrections during the production process.
We also propose more traditional post-processing techniques, such a fast approximate anti-aliasing, or a ray-marched screen-space ambient occlusion method [Sourimant et al. 2010] that — contrary to concurrent approaches in other viewports — is designed to match MentalRay's parameterization and results, making it a usable ambient occlusion previsualization, or even a fast and low-cost ambient occlusion rendering pass.

[1]Mental Ray ©NVidia — www.mentalimages.com
[2]RenderMan ©Pixar — www.renderman.pixar.com
[3]Maya ©Autodesk — www.autodesk.com

References

ESTEVE, J., PORTSMOUTH, J., GAUTRON, P., PRUNIER, J.-C., MARVIE, J.-E., AND DELALANDRE, C. 2012. Bringing transmittance function maps to the screen: Wrath of the titans and prometheus. In *Proceedings of the Digital Production Symposium*, DigiPro'12.

GAUTRON, P., DELALANDRE, C., MARVIE, J.-E., AND LECOCQ, P. 2013. Boundary-aware extinction mapping. *Computer Graphics Forum 32*, 7, 305–314.

MARVIE, J.-E., AND BOUATOUCH, K. 2003. Remote rendering of massively textured 3d scenes through progressive texture maps. In *proceedings of the third IASTED conference on Visualisation Imaging and Image Processing*, 756–761.

MARVIE, J.-E., HIRTZLIN, P., AND GAUTRON, P. 2013. Many-core event evaluation. In *Proceedings of the 18th International Conference on 3D Web Technology*, Web3D'13.

SOURIMANT, G., GAUTRON, P., AND MARVIE, J.-E. 2010. Poisson disk ray-marched ambient occlusion. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, I3D Posters*.